

1. (Currently Amended) In a database system, a method of maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range comprising:

creating at least one new bucket in response to a query on the database, [[wherein the]] each new bucket having [[has]] bucket boundaries corresponding to a range of tuple attribute values returned by the query and a bucket frequency corresponding to a number of tuples returned by the query; ~~and wherein each new bucket is contained within at least one existing bucket and wherein the new bucket becomes a child bucket that has a child bucket frequency and the existing bucket becomes a parent bucket that has a parent bucket frequency~~

establishing a logical relationship between the new bucket and an existing bucket such that the existing bucket is a parent bucket of the new bucket;

storing the self-tuning histogram that includes the new bucket in memory; and

wherein bucket boundaries of each new bucket fall within bucket boundaries of the parent bucket of the new bucket.

2. (Original) The method of claim 1 wherein each bucket has a rectangular shape.

3. (Canceled)

4. (Currently Amended) The method of claim 1 wherein ~~a total number of buckets is limited to a predetermined budget~~ a child bucket forms a hole in the parent bucket of the child bucket.

5. (Currently Amended) The method of claim 1, further ~~[[4]]~~ comprising merging buckets based on a merge criterion when the total number of buckets exceeds a ~~[[the]]~~ predetermined budget.

6. (Original) The method of claim 5 wherein the merge criterion is a similar bucket density, wherein bucket density is based on the bucket frequency divided by the bucket volume.

7. (Currently Amended) The method of claim 1 further comprising shrinking the boundaries of ~~[[each]]~~ a new bucket if the boundaries of the new bucket intersect any existing bucket boundaries.

8. (Original) The method of claim 1 wherein the frequency of the parent bucket is diminished by the frequency of the child bucket.

9. (Currently Amended) In a database system, a method of maintaining a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attributes that fall in the bucket boundary range, the method comprising:

- a) examining the results of a query executed on the database;
- b) creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query;
- c) modifying the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket; [[and]]
- d) creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole; and
- e) storing the modified self-tuning histogram in one or more computer-readable media.

10. (Original) The method of claim 9 wherein each bucket has a rectangular shape.

11. (Canceled)

12. (Original) The method of claim 9 wherein a total number of buckets is limited to a predetermined budget.

13. (Previously Amended) The method of claim 12 comprising merging buckets based on a merge criterion when the total number of buckets exceeds the predetermined budget.

14. (Original) The method of claim 13 wherein the merge criterion is a similar bucket density, wherein bucket density is based on the bucket frequency divided by the bucket volume.

15. (Original) The method claim 9 wherein the frequency of the parent bucket is diminished by the frequency of the child bucket.

16. (Currently Amended) One or more [[A]] computer readable media [[medium]] having [[computer]] executable instructions that, when executed, implement [[for performing steps]] a method for maintaining a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range, the steps comprising:

- a) examining the results of a query executed on the database;
- b) creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute

values returned by the query and a frequency corresponding to a number of tuples returned by the query;

c) modifying the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket; and

d) creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole; and

e) storing the modified self-tuning histogram in one or more computer-readable media.

17. (Currently Amended) The one or more computer readable media [[medium]] of claim 16 wherein each bucket has a rectangular shape.

18. (Canceled)

19. (Currently Amended) The one or more computer readable media [[medium]] of claim 16 wherein the method [[steps comprise]] further comprises merging buckets having a similar bucket density when the total number of buckets exceeds a predetermined budget.

20. (Previously Amended) An apparatus for maintaining a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume,

and a bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range comprising:

- a) means for examining the results of a query executed on the database;
- b) means for creating at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query;
- c) means for modifying the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket; and
- d) means for creating a new child bucket that has a child frequency in the histogram corresponding to each modified hole.

21. (Canceled)

22. (Previously Amended) An apparatus that maintains a self-tuning histogram having a plurality of existing parent buckets arranged in a hierarchical manner and defined by at least two bucket boundaries that represent a range of attribute values, a bucket volume, and bucket frequency that corresponds to a number of tuples having attribute values that fall in the bucket boundary range comprising:

- a) a memory device that stores a database comprising multiple data records;

- b) a computer having one or more processing units that execute a stored computer program, said computer including a rapid access memory store; and
- c) an interface that couples the memory device that stores the database to the computer to allow records to be retrieved from the database; wherein
- d) the stored program has components including i) a component that examines the results of a query executed on the database; ii) a component that creates at least one candidate hole in the histogram based on the results of the query such that the candidate hole has boundaries corresponding to a range of attribute values returned by the query and a frequency corresponding to a number of tuples returned by the query; iii) a component that modifies the boundaries of each candidate hole such that the boundaries of the modified hole are completely contained within the boundaries of at least one existing parent bucket and do not partially intersect the boundaries of any existing bucket; and iv) a component that creates a new child bucket that has a child frequency in the histogram corresponding to each modified hole.

23. (Canceled)

24. (Currently Amended) For use with a database system, a histogram tuning system comprising:

a component that receives a [[bucket from the]] histogram having at least a parent bucket; and

a tuning component that iteratively populates the parent bucket with a child bucket, as a function of query results, wherein the child bucket is completely contained within the parent bucket.

25. (Currently Amended) The histogram tuning system of claim 24 wherein the tuning component populates the parent bucket with a child bucket that has boundaries corresponding to a range of attribute values present in the query results and a child bucket frequency corresponding to a number of tuples present in the query results.

26. (Previously Presented) The histogram tuning system of claim 24 comprising a merging component that merges buckets based on a merge criteria.

27. (Previously Presented) The histogram tuning system of claim 26 wherein the merge criterion is a similar bucket density.

28. (Previously Presented) The histogram tuning system of claim 24 wherein the tuning component shrinks the boundaries of the child bucket if the child bucket boundaries intersect any other bucket boundaries.

29. (Currently Amended) The histogram tuning system of claim 25 wherein the parent bucket has a parent bucket frequency and the parent bucket frequency is diminished by the child bucket frequency.



30. (Previously Presented) A database histogram tuning system comprising:

means for receiving a bucket from a histogram; and

means for iteratively populating the bucket with a child bucket, as a function of query results, such that the child bucket is fully contained within the received bucket.